



HCIIF™ WHITE PAPER

When Can Your Organisation Rely on Its Information?

An introduction to the Human-Centric Information Integrity Framework for decisions, information sharing and responsible AI.

Deepak Sadasivan

Creator of HCIIF™

First Edition · October 2026

hciif.com

Executive summary

Organisations need to know when information is suitable for the decision or activity in front of them. Information may be accurate and secure, yet its assumptions, limitations or original purpose may no longer match the proposed use. Sharing, transformation and changing conditions can weaken the basis for earlier reliance.

The Human-Centric Information Integrity Framework (HCIIF™) is an Enterprise Information Integrity Management System. It provides a structured approach to establishing, governing and maintaining confidence in information, based on evidence. Its scope includes records, reports, assessments and information used or produced by AI. [Enterprise Edition, 1, Chapters 1–4]

How HCIIF works

HCIIF identifies the information that matters, assesses the evidence for its intended use, explains the confidence justified by that evidence and makes the relevant context available. Before relying on information, the organisation checks that the assessment applies, confidence is sufficient, the use is permitted and the decision sits with the right authority. [Enterprise Edition, 1, Chapters 5–10 and 13]

Confidence remains linked to evidence, purpose, limitations and conditions. It does not itself grant permission or authority. When those conditions change, the organisation reviews affected uses and communicates the action required.

Why this matters

The intended benefit is clearer reasoning at the point of use: people can understand why reliance is justified, what remains uncertain and when to seek more evidence or restrict an activity. AI-generated information also needs assessment for its intended use; confidence in its sources does not automatically establish confidence in its output.

Established guidance recognises the importance of fitness for purpose, provenance and context-sensitive AI assessment. [5–8] HCIIF connects these concerns through an explicit basis for reliance and accountable decisions.

Conclusion and practical next step

HCIIF's practical benefits need evaluation in real settings. Begin with one important activity and compare decision traceability, responses to change and effort with a documented baseline. A focused pilot or research collaboration can test whether the approach is useful and proportionate. The examples in this paper illustrate possible applications; they are not validated pilot results.

Contents and supporting documents

Choose your starting point

For a strategic overview, read the executive summary on page 2 and the conclusions on page 11. For practical detail, follow the challenge, framework, use, change and AI sections on pages 4–9, then the examples on page 10.

Section	Page
Executive summary	2
The information integrity challenge	4
Research and context	5
The HCIIF architecture	6
From evidence to responsible use	7
Maintaining confidence as conditions change	8
Information integrity before, through and after AI	9
Two illustrative applications	10
Adoption and evaluation	11
References and author	12

Supporting HCIIF documents

The following documents provide the detailed reference and implementation material behind this paper.

Enterprise Edition — defines the formal architecture, principles, terms and relationships. It is the authoritative HCIIF reference. [1]

Practitioner Guide — explains how organisations can implement and adopt the framework. [2]

Practitioner Toolkit — provides connected working records to support implementation. [3]

Technical Workflow — illustrates the operational and technical process. [4]

The information integrity challenge

Information can reach a decision-maker without the context needed to judge its use. A dashboard may omit assumptions. A shared assessment may lose its restrictions. A summary may retain a conclusion while leaving out disputed findings or dated evidence. The information is still available, but the reasons and conditions for relying on it are harder to see.

A team may then reuse an assessment for a more consequential decision, share information beyond its permitted scope or continue acting on evidence that needs review. Data quality, security, records management, legal, privacy and risk controls provide essential evidence. HCIIF connects that evidence to a practical question: can we rely on this information for this use, under these conditions? [Enterprise Edition, 1, Chapters 1–4]




Figure 1 The same supplier assessment may support routine planning but need further evidence for a critical service approval. Illustrative example.

The Confidence Gap

A Confidence Gap can arise when assumed confidence exceeds what the evidence justifies, or when justified confidence falls below Required Confidence — the confidence needed for the intended use. Making these comparisons explicit helps people recognise where reliance needs further consideration. [Enterprise Edition, 1, Chapters 1 and 5]

In the supplier example, an assessment used for planning may leave questions about service resilience unanswered. Those questions become important when approving a critical dependency. The proposed use and its consequences determine whether further evidence or safeguards are needed.

The relevant authority can seek more evidence, limit the use, add safeguards, accept defined exposure where permitted or stop the activity. An official source or a confident presentation is one part of the context; the decision still needs evidence appropriate to its purpose.

Research and context

Established guidance provides a basis for understanding information integrity. Three themes are particularly relevant: information must fit its intended purpose, its source and history should be understood, and its suitability needs review as conditions change.

Quality depends on what information is used for

The UK Government Data Quality Framework places user needs and fitness for purpose at the centre of quality management. It recognises that the same data may serve some purposes better than others and calls for quality management throughout its lifecycle. This supports checking suitability when information is reused for a different activity. [5]

Provenance helps explain the evidence

The World Wide Web Consortium (W3C) PROV Data Model describes the information, activities and people involved in producing a record or result. This provenance helps assess quality, reliability and trustworthiness. Its 2013 Recommendation provides the established provenance concepts referenced here. [6]

AI assessment needs context and continued review

The US National Institute of Standards and Technology (NIST) AI Risk Management Framework 1.0 links trustworthiness to context, human judgement and evaluation. Its Generative AI Profile identifies confidently presented false or erroneous content as a risk and includes measures to verify sources and citations. These sources support evaluating AI-generated information for its intended use. [7–8]

What HCIIF adds

HCIIF brings evidence from existing quality, security, privacy and governance controls into a connected approach to managing reliance. It assesses justified confidence, explains the Confidence Basis, communicates relevant context and checks the conditions and authority for use. Review and notification arrangements address changes that may affect earlier reliance. [Enterprise Edition, 1, Chapters 3–7, 10 and 13]

What practical evaluation should establish

The external sources explain why these concerns matter; evidence of HCIIF's practical value must come from its application. Evaluation should test whether people can understand and use the confidence context, whether affected users respond appropriately to change, and whether the benefit justifies the effort. These findings can guide improvement and decisions about wider adoption.

The HCIIF architecture

HCIIF connects information, evidence, confidence and the conditions for use within an organisational management system. Leadership, planning, operation, assurance, Management Review and improvement support these relationships. Adoption can begin with one defined service, programme or decision. The figure shows a simplified operational view of the framework. [Enterprise Edition, 1, Chapters 3–4]




Figure 2 A simplified view of how HCIIF supports reliance. Assurance and improvement inform each stage as evidence, uses and conditions change.

EIA = Enterprise Information Asset: information that needs deliberate integrity management.

ECAM = Enterprise Confidence Assessment Methodology: the method for assessing evidence and justified confidence.

IIP = Information Integrity Passport: the record or view that communicates relevant integrity context.

Identify the information that matters

Define an EIA around information that matters to a decision or activity. Its scope should make assessment, ownership and action practical. It may draw on several records or systems, rather than match a single file or application. [Enterprise Edition, 1, Chapter 8]

Connect the assessment to its use

ECAM supports an Enterprise Confidence position for a defined purpose. The Confidence Basis explains the evidence, assumptions, limitations and conditions that justify that position. The IIP makes the relevant context visible and traceable to users and systems. Protected evidence can remain restricted while appropriate assurance and limitations are communicated. [Enterprise Edition, 1, Chapters 5–6]

Existing records and systems can carry this context if they preserve its meaning, links and history. The aim is to make the basis for reliance available where information is used, with assessment and governance proportionate to the consequences. [Practitioner Guide, 2, Chapters 3 and 17]

From evidence to responsible use

Start with the intended use and its possible consequences. Required Confidence is the confidence needed for that use. ECAM examines the supporting evidence through lenses such as suitability, relevance, accuracy, completeness, timeliness, reliability, consistency, provenance, authority and Information Rights. The weight given to each lens depends on the context. [Enterprise Edition, 1, Chapter 5]

Explain Enterprise Confidence

The baseline Enterprise Confidence levels are High, Moderate, Low and Insufficient. They communicate the strength of the current justification for reliance in context. Each level stays attached to its Confidence Basis, purpose, assessment date, limitations and review conditions. A level is not a probability of correctness: High does not guarantee suitability for every use, and Low does not necessarily mean the information is false. [Enterprise Edition, 1, Chapter 5]

Before using the information	What must be understood
Does the assessment apply here?	Current version, intended purpose, recipient, context and material transformation.
Is confidence sufficient?	Enterprise Confidence and its Confidence Basis compared with Required Confidence for the use.
Is this use permitted?	Information Rights, applicable Caveats and whether their conditions are satisfied.
Who may decide and act?	Authority to act, accept reliance below Required Confidence where permitted, or escalate.

Permission and accountable authority

Information Rights define what may be done with information. Caveats set the conditions or restrictions on those rights. A well-supported confidence assessment can therefore coexist with a restriction on access, sharing or use. The organisation must check both the evidence and the applicable permissions. [Enterprise Edition, 1, Chapters 5 and 10]

Confidence Authority governs the assessment, approval or challenge of Enterprise Confidence within an agreed scope. Decision Authority governs action. Acceptance Authority governs acceptance of reliance despite a known Confidence Gap. These responsibilities can sit within existing roles; they do not automatically require new job titles. [Enterprise Edition, 1, Chapter 13]

Where reliance below Required Confidence is permitted and authorised, record the decision, reasons, scope, safeguards and review or expiry conditions. Acceptance records a decision to proceed; it leaves the Confidence Gap visible and respects independent restrictions. [Enterprise Edition, 1, Chapter 13]

Maintaining confidence as conditions change

Confidence remains linked to a defined purpose and assessment date. New evidence, corrections, transformations or changed conditions can affect the basis for reliance. A Confidence Event is an occurrence, observation or change that may materially affect that basis and warrants consideration. First assess its significance: an event can strengthen, leave unchanged or weaken justified confidence.

[Enterprise Edition, 1, Chapters 5 and 7]

Confidence Drift is gradual movement in the justified basis for Enterprise Confidence as changes accumulate. It can be positive or negative. For example, one late feed may be an event; repeated late feeds may reveal negative Confidence Drift. Monitoring both events and trends helps identify when review is needed. [Enterprise Edition, 1, Chapter 7]




Figure 3 Review the significance of change before deciding the response. Immaterial events may be recorded and monitored without full reassessment; material events require evidence review and reassessment where necessary.

IIP = Information Integrity Passport. Preserve version history and communicate relevant changes to affected users and systems.

Respond through Confidence Recall

A Confidence Recall is a governed notification that a previous confidence position, information version or basis for reliance needs reconsideration following a material development. It does not necessarily mean the information was false. Identify affected versions, uses, recipients and dependencies, preserve the earlier record and communicate the review or restrictions required. Any decision to resume suspended reliance needs an appropriate current basis and authority. [Enterprise Edition, 1, Chapter 9; Practitioner Guide, 2, Chapters 10–11]

Manage Confidence Debt and recurring weaknesses

Confidence Debt is accumulated exposure when known weaknesses remain unresolved while reliance continues. Material debt should have an owner, a defined response, a review date and appropriate acceptance or escalation. Assurance Intelligence identifies patterns in integrity evidence; Management Review uses them to prioritise improvement. [Enterprise Edition, 1, Chapters 7 and 11]

Information integrity before, through and after AI

AI can summarise, combine and transform information, changing its meaning or the basis for its use. Confidence in a source does not automatically establish Enterprise Confidence in a generated result. Where transformation creates materially distinct information, assess the result for its intended use, with effort proportionate to its consequences and uncertainty. [Enterprise Edition, 1, Chapters 17–18 and 20; Practitioner Guide, 2, Chapters 6 and 18]




Figure 4 Preserve relevant source context through transformation and assess the resulting information before reliance. Review should reflect the intended use and its consequences.

Illustrative reporting example

An AI assistant summarises operational reports for a manager. Some observations are dated, some findings are disputed and some reports restrict onward sharing. If the summary presents every finding as current and agreed, it changes the apparent basis for a decision. Review should check that important qualifications and use conditions remain visible, that claims can be traced to supporting evidence and that the result is suitable for the manager's intended action.

Meaningful human accountability

Human-centricity places accountability for purpose, evidence, use conditions and action with people and their governed organisational arrangements. Depending on the activity, this may involve direct review, supervised operation or governed automation. People need the skills, time, authority and practical means to challenge, intervene and recover. Oversight should fit the use and its consequences, rather than require approval of every automated event. [Practitioner Guide, 2, Chapter 18]

Technology enablement

The Enterprise Confidence Enablement Layer (CEL) is a logical integration layer through which existing technology can carry and expose confidence context. It can support evidence collection, visible use conditions, change detection and exception routing. CEL requires no particular platform, and automated outputs contribute evidence rather than independently create Enterprise Confidence. HCIF's information integrity focus complements wider model risk, safety, security and legal governance. [Enterprise Edition, 1, Chapters 3 and 12; Technical Workflow, 4, page 1]

Two illustrative applications

These hypothetical scenarios show how HCIIF could support practical decisions. Each identifies a problem, a possible application and questions for evaluation. They do not report completed deployments or measured benefits.

Supplier information reused for a new decision

Situation. A supplier assessment supports routine planning but notes that service recovery has not been independently tested. A team later uses a summary of that assessment when considering a critical service dependency. The summary omits the limitation, making the evidence appear stronger than it is.

Application. Define the proposed decision and Required Confidence, confirm the assessment version and dependencies, and review the evidence through ECAM. The Confidence Basis would keep the recovery limitation visible; the IIP would communicate the confidence position and use conditions. The relevant authority could seek recovery evidence, require safeguards or limit the proposed reliance. If new findings change the basis, affected uses would be identified for review. [Enterprise Edition, 1, Chapters 5–9 and 13]

Evaluation. Test whether reviewers identify the missing limitation, trace the decision to evidence and authority, and respond to changed findings. Compare results and effort with the agreed baseline.

Sharing a sensitive assessment with a partner

Situation. An organisation plans to share a sensitive assessment with a partner for a defined operational purpose. The partner needs enough evidence to judge its use, while some underlying sources must remain protected. Sharing the assessment must also preserve restrictions on onward disclosure.

Application. Before transfer, establish authority, check Caveats and assess Transfer Suitability for the recipient and receiving environment. This considers the proposed purpose, sensitivity, restrictions, recipient capability and consequences of sharing. The IIP could communicate a restricted assurance summary and relevant limitations. The recipient would assess whether the confidence position applies and meets Required Confidence for its own use. An agreed notification route would support review if the original basis changes. [Enterprise Edition, 1, Chapters 6 and 10; Practitioner Guide, 2, Chapter 20]

Evaluation. Test whether recipients understand and observe the use conditions, whether the assurance summary supports their assessment, and whether change notifications reach affected users and prompt traceable responses.

Judge the decisions and responses

In both scenarios, evaluation should examine whether the integrity context changes how information is used. A completed record is evidence of activity. The stronger test is whether a decision or response is justified, permitted, accountable and proportionate to the consequence.

Adoption and evaluation

Start with one decision or activity affected by lost context, unclear confidence or changing conditions. Agree the intended use, scope, sponsor, information owner and relevant authorities. Define what the pilot will test and how results will be assessed. Use existing systems where they can preserve context, relationships and history. [Practitioner Guide, 2, Chapters 3, 15 and 17; Practitioner Toolkit, 3, T20]

Stage	Output to review
Define the need	Intended use, scope, consequences, owners and agreed measures.
Map and assess	Information and dependencies; Required Confidence; findings and reasons.
Test use and change	An IIP used in a decision; a controlled change exercise and recorded responses.
Evaluate and decide	Results and effort compared with the starting position; a decision to improve, expand or stop.

Measure value and effort

Record how the activity works before the pilot, or agree another suitable comparison. Check whether sampled decisions can be traced to their evidence and authority, how quickly affected dependencies are found after a change, and whether notifications lead to appropriate action. Measure the effort needed to maintain records as well as any benefit. Report the evidence, review period and limits of the findings.

Key takeaways

Confidence applies to a defined purpose and time, supported by evidence and a clear account of its limitations. Permission and decision authority require separate consideration. Changes may require earlier reliance to be reviewed. Information produced through AI needs assessment in its own context, alongside the information used to produce it.

Conclusion and next step

HCIIF provides a structured way to keep information, evidence, confidence and accountability connected. A focused pilot can test whether this helps people make and review decisions, and whether the effort is proportionate. Organisations and research partners interested in collaboration can bring a specific information challenge, a proposed scope and the questions they want to test.

Visit hciif.com to explore a pilot or collaboration.

References and author

References [1–4] are the HCIIF source documents. References [5–8] provide independent context; they do not constitute validation or endorsement of HCIIF.

[1] Sadasivan, D. HCIIF™ Enterprise Edition.

Defines the architecture, principles, terms and relationships.

[2] Sadasivan, D. HCIIF™ Practitioner Guide.

Explains implementation and adoption.

[3] Sadasivan, D. HCIIF™ Practitioner Toolkit.

Provides connected implementation records.

[4] Sadasivan, D. HCIIF™ Technical Workflow.

Illustrates the operational and technical process.

[5] UK Government. The Government Data Quality Framework (2020).

[Read the source](#)

[6] W3C. PROV-DM: The PROV Data Model. W3C Recommendation (30 April 2013).

[Read the source](#)

[7] NIST. Artificial Intelligence Risk Management Framework 1.0. NIST AI 100-1 (2023).

[Read the source](#)

[8] NIST. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. NIST AI 600-1 (2024).

[Read the source](#)

About the author

Deepak Sadasivan is a senior consultant and strategic advisor with more than 20 years of operational and consulting experience across sensitive and highly regulated environments. His work in information governance, international sharing and enterprise transformation informed HCIIF. [1, About the Author]

HCIIF brings existing governance and assurance evidence together to help people understand when information can be relied on, under what conditions and who is accountable for its use.